

CLAIMS

1. A method for storing data in a data storage system that includes primary and secondary storage subsystems, including respective first and second volatile cache memories and
5 respective first and second non-volatile storage media, the method comprising:
 - receiving the data at the primary storage subsystem from a host processor;
 - 10 writing the data to the first volatile cache memory in the primary storage subsystem;
 - copying the data from the primary storage subsystem to the secondary storage subsystem;
 - writing the copied data to the second volatile cache memory in the secondary storage subsystem;
 - 15 returning an acknowledgment from the secondary storage subsystem to the primary storage subsystem responsively to writing the copied data to the second volatile cache memory and prior to saving the data in the second non-volatile storage media;
 - 20 signaling the host processor that the data have been stored in the data storage system responsively to the acknowledgment from the secondary storage subsystem; and
 - transferring the data in the primary and secondary storage subsystems from the first and second volatile cache
25 memories to the first and second non-volatile storage media, respectively.

2. The method according to claim 1, wherein copying the data comprises transmitting the data between mutually-remote sites over a communication link between the sites.

5 3. The method according to claim 1, wherein the second volatile cache memory and the second non-volatile storage media are located in mutually-remote sites, and wherein transferring the data comprises transmitting the data from the second volatile cache memory to the second non-volatile
10 storage media over a communication link between the sites.

4. The method according to claim 1, wherein copying the data comprises creating a mirror on the secondary storage subsystem of the data received by the primary storage
15 subsystem.

5. The method according to claim 4, and comprising, upon occurrence of a failure in the primary storage subsystem, configuring the secondary storage subsystem to serve as the
20 primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

6. The method according to claim 4, wherein the second volatile cache memory is located in a site remote from the
25 primary storage subsystem and is operated by a service provider other than an owner of the primary storage subsystem, and wherein transferring the data from the second

volatile cache memory to the second non-volatile storage media comprises processing the data on a fee-per-service basis.

7. The method according to claim 1, wherein transferring
5 the data comprises sending a message from the secondary storage subsystem to the primary storage subsystem indicating addresses of the data that have been transferred to the second non-volatile storage media, and wherein the method further comprises creating a record on the primary
10 storage subsystem of the addresses of the data copied to the secondary storage subsystem, and updating the record in response to the message.

8. The method according to claim 7, and further
15 comprising, upon recovery of the system from a failure of the secondary storage subsystem, conveying, responsively to the record, a portion of the data from the primary storage subsystem to be stored on the secondary storage subsystem.

20 9. The method according to claim 8, wherein updating the record comprises removing from the record the addresses of the data that have been transferred to the second non-volatile storage media.

25 10. The method according to claim 7, wherein creating the record comprises marking respective bits in a bitmap corresponding to addresses of the data copied to the

secondary storage subsystem, and wherein updating the record comprises clearing the respective bits.

11. The method according to claim 7, wherein transferring
5 the data comprises transferring the data in a range of the
addresses from the second volatile cache memory to the
second non-volatile storage media, and wherein sending the
message comprises informing the primary storage subsystem
that the data in the range have been transferred, so that
10 the primary storage subsystem updates the record with
respect to the range.

12. The method according to claim 11, wherein transferring
the data in the range comprises destaging the range of the
15 addresses.

13. The method according to claim 11, wherein creating the
record comprises maintaining a mirror of the record on the
secondary storage subsystem, and wherein transferring the
20 data in the range comprises selecting the data to be
transferred from the second volatile cache memory to the
second non-volatile storage media responsively to the
mirror.

25 14. The method according to claim 1, wherein transferring
the data comprises sending a message from the primary
storage subsystem to the secondary storage subsystem

indicating addresses of the data that have been transferred to the first non-volatile storage media, and wherein the method further comprises creating a record on the secondary storage subsystem of the addresses of the data copied to the secondary storage subsystem, and updating the record in response to the message.

15. The method according to claim 14, and further comprising, upon recovery of the system from a failure of the primary storage subsystem, conveying, responsively to the record, a portion of the data from the secondary storage subsystem to be stored on the primary storage subsystem.

16. The method according to claim 15, wherein updating the record comprises removing from the record the addresses of the data that have been transferred to the primary non-volatile storage media.

17. The method according to claim 14, wherein creating the record comprises marking respective bits in a bitmap corresponding to addresses of the data copied to the secondary storage subsystem, and wherein updating the record comprises clearing the respective bits.

18. The method according to claim 14, wherein transferring the data comprises transferring the data in a range of the addresses from the first volatile cache memory to the first

non-volatile storage media, and wherein sending the message comprises informing the secondary storage subsystem that the data in the range have been transferred, so that the secondary storage subsystem updates the record with respect
5 to the range.

19. The method according to claim 18, wherein transferring the data in the range comprises destaging the range of the addresses.

10

20. The method according to claim 18, wherein creating the record comprises maintaining a mirror of the record on the primary storage subsystem, and wherein transferring the data in the range comprises selecting the data to be transferred
15 from the first volatile cache memory to the first non-volatile storage media responsively to the mirror.

21. A data storage system, comprising:

20 a primary storage subsystem, which comprises a first volatile cache memory and first non-volatile storage media;
and

a secondary storage subsystem, which comprises a second volatile cache memory and second non-volatile storage media,

25 wherein the primary storage subsystem is arranged to receive data from a host processor, to write the data to the first volatile cache memory, to copy the data to the secondary storage subsystem, and to transfer the data from

the first volatile cache memory to the first non-volatile storage media, and

wherein the secondary storage subsystem is arranged to receive and write the copied data to the second volatile cache memory, to transfer the data from the first volatile
5 cache memory to the first non-volatile storage media, and to return an acknowledgment to the primary storage subsystem responsively to writing the copied data to the second volatile cache memory and prior to transferring the data to
10 the second non-volatile storage media,

wherein the primary storage subsystem is arranged to signal the host processor that the data have been stored in the data storage system responsively to the acknowledgment from the secondary storage subsystem.

15

22. The system according to claim 21, wherein the primary and secondary storage subsystems are located at mutually-remote sites, and are coupled to communicate over a communication link between the sites.

20

23. The system according to claim 21, wherein the second volatile cache memory and the second non-volatile storage media are located in mutually-remote sites, and wherein the secondary storage subsystem is arranged to transfer the data
25 from the second volatile cache memory to the second non-volatile storage media over a communication link between the sites.

24. The system according to claim 21, wherein the secondary storage subsystem is arranged to mirror the data held by the primary storage subsystem.

5 25. The system according to claim 24, wherein the secondary storage subsystem is configurable, upon occurrence of a failure in the primary storage subsystem, to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the data storage system.

10

26. The system according to claim 21, wherein the second volatile cache memory is located in a site remote from the primary storage subsystem and is operated on a fee-per-service basis by a service provider other than an owner of
15 the primary storage subsystem.

27. The system according to claim 21, wherein the secondary storage subsystem is arranged to send a message to the primary storage subsystem indicating addresses of the data
20 that have been transferred to the second non-volatile storage media, and wherein the primary storage subsystem is arranged to create a record of the addresses of the data copied to the secondary storage subsystem, and to update the record in response to the message.

25

28. The system according to claim 27, wherein the primary storage subsystem is arranged, upon recovery of the system

from a failure of the secondary storage subsystem, to convey, responsively to the record, a portion of the data from the primary storage subsystem to be stored on the secondary storage subsystem.

5

29. The system according to claim 28, wherein the primary storage subsystem is arranged to remove from the record the addresses of the data that have been transferred to the second non-volatile storage media.

10

30. The system according to claim 27, wherein the primary storage subsystem is arranged to create the record by marking respective bits in a bitmap corresponding to addresses of the data copied to the secondary storage subsystem, and to update the record by clearing the respective bits in response to the message.

15

31. The system according to claim 27, wherein the secondary storage subsystem is arranged to transfer the data over a range of the addresses from the second volatile cache memory to the second non-volatile storage media, and to indicate in the message that the data in the range have been transferred, so that the primary storage subsystem updates the record with respect to the range.

25

32. The system according to claim 31, wherein the secondary storage subsystem is arranged to transfer the data by destaging the range of the addresses.

5 33. The system according to claim 31, wherein the secondary storage subsystem is arranged to maintain a mirror of the record, and to select the data to be transferred from the second volatile cache memory to the second non-volatile storage media responsively to the mirror.

10

34. The system according to claim 21, wherein the primary storage subsystem is arranged to send a message to the secondary storage subsystem indicating addresses of the data that have been transferred to the first non-volatile storage
15 media, and wherein the secondary storage subsystem is arranged to create a record of the addresses of the data copied to the secondary storage subsystem, and to update the record in response to the message.

20 35. The system according to claim 34, wherein the secondary storage subsystem is arranged, upon recovery of the system from a failure of the primary storage subsystem, to convey, responsively to the record, a portion of the data from the secondary storage subsystem to be stored on the primary
25 storage subsystem.

36. The system according to claim 35, wherein the secondary storage subsystem is arranged to remove from the record the addresses of the data that have been transferred to the first non-volatile storage media.

5

37. The system according to claim 34, wherein the secondary storage subsystem is arranged to create the record by marking respective bits in a bitmap corresponding to addresses of the data copied to the secondary storage subsystem, and to update the record by clearing the
10 respective bits in response to the message.

38. The system according to claim 34, wherein the primary storage subsystem is arranged to transfer the data over a
15 range of the addresses from the first volatile cache memory to the first non-volatile storage media, and to indicate in the message that the data in the range have been transferred, so that the secondary storage subsystem updates the record with respect to the range.

20

39. The system according to claim 38, wherein the primary storage subsystem is arranged to transfer the data by destaging the range of the addresses.

25 40. The system according to claim 38, wherein the primary storage subsystem is arranged to maintain a mirror of the record, and to select the data to be transferred from the

first volatile cache memory to the first non-volatile storage media responsively to the mirror.

41. A computer software product for use in a data storage
5 system including primary and secondary storage subsystems,
which include respective first and second control units,
respective first and second volatile cache memories, and
respective first and second non-volatile storage media, the
product comprising a computer-readable medium in which
10 program instructions are stored, which instructions, when
read by the first and second control units, cause the first
control unit to receive data from a host processor, to write
the data to the first volatile cache memory, to copy the
data to the secondary storage subsystem, and to transfer the
15 data from the first volatile cache memory to the first non-
volatile storage media, and cause the second control unit to
receive and write the copied data to the second volatile
cache memory, to transfer the data from the first volatile
cache memory to the first non-volatile storage media, and
20 prior to transferring the data to the second non-volatile
storage media, to return an acknowledgment to the primary
storage subsystem responsively to writing the copied data to
the second volatile cache memory, wherein the instructions
further cause the first control unit to signal the host
25 processor that the data have been stored in the data storage
system responsively to the acknowledgment from the secondary
storage subsystem.

42. The product according to claim 41, wherein the primary and secondary storage subsystems are located at mutually-remote sites, and wherein the instructions cause the first and second control units to communicate over a communication
5 link between the sites.

43. The product according to claim 41, wherein the second volatile cache memory and the second non-volatile storage media are located in mutually-remote sites, and wherein the
10 instructions cause the second control unit to transfer the data from the second volatile cache memory to the second non-volatile storage media over a communication link between the sites.

44. The product according to claim 41, wherein the
15 instructions cause the first and second control units to mirror the data held by the primary storage subsystem on the secondary storage subsystem.

45. The product according to claim 44, wherein the
20 instructions cause the secondary storage subsystem, upon occurrence of a failure in the primary storage subsystem, to serve as the primary storage subsystem so as to receive further data from the host processor to be stored by the
25 data storage system.

46. The product according to claim 44, wherein the second volatile cache memory is located in a site remote from the primary storage subsystem and are operated by a service provider other than an owner of the primary storage subsystem, and wherein the instructions cause the second control unit to transfer the data from the second volatile cache memory to the second non-volatile storage media on a fee-per-service basis.

47. The product according to claim 41, wherein the instructions cause the second control unit to send a message to the primary storage subsystem indicating addresses of the data that have been transferred to the second non-volatile storage media, and further cause the first control unit to create a record of the addresses of the data copied to the secondary storage subsystem, and to update the record in response to the message.

48. The product according to claim 47, wherein the instructions cause the first control unit, upon recovery of the system from a failure of the secondary storage subsystem, to convey, responsively to the record, a portion of the data from the primary storage subsystem to be stored on the secondary storage subsystem.

49. The product according to claim 48, wherein the instructions cause the first control unit to remove from the

record the addresses of the data that have been transferred to the second non-volatile storage media.

50. The product according to claim 47, wherein the
5 instructions cause the first control unit to create the record by marking respective bits in a bitmap corresponding to addresses of the data copied to the secondary storage subsystem, and to update the record by clearing the respective bits in response to the message.

10

51. The product according to claim 47, wherein the instructions cause the second control unit to transfer the data over a range of the addresses from the second volatile cache memory to the second non-volatile storage media, and
15 to indicate in the message that the data in the range have been transferred, so that the first control unit updates the record with respect to the range.

52. The product according to claim 51, wherein the
20 instructions cause the second control unit to transfer the data by destaging the range of the addresses.

53. The product according to claim 51, wherein the instructions cause the second control unit to maintain a
25 mirror of the record, and to select the data to be transferred from the second volatile cache memory to the

second non-volatile storage media responsively to the mirror.

54. The product according to claim 41, wherein the
5 instructions cause the first control unit to send a message
to the secondary storage subsystem indicating addresses of
the data that have been transferred to the first non-
volatile storage media, and further cause the second control
unit to create a record of the addresses of the data copied
10 to the secondary storage subsystem, and to update the record
in response to the message.

55. The product according to claim 54, wherein the
instructions cause the second control unit, upon recovery of
15 the system from a failure of the primary storage subsystem,
to convey, responsively to the record, a portion of the data
from the secondary storage subsystem to be stored on the
primary storage subsystem.

20 56. The product according to claim 55, wherein the
instructions cause the second control unit to remove from
the record the addresses of the data that have been
transferred to the first non-volatile storage media.

25 57. The product according to claim 54, wherein the
instructions cause the second control unit to create the
record by marking respective bits in a bitmap corresponding

to addresses of the data copied to the secondary storage subsystem, and to update the record by clearing the respective bits in response to the message.

5 58. The product according to claim 54, wherein the instructions cause the first control unit to transfer the data over a range of the addresses from the first volatile cache memory to the first non-volatile storage media, and to indicate in the message that the data in the range have been
10 transferred, so that the secondary storage subsystem updates the record with respect to the range.

59. The product according to claim 58, wherein the instructions cause the first control unit to transfer the
15 data by destaging the range of the addresses.

60. The product according to claim 58, wherein the instructions cause the first control unit to maintain a mirror of the record, and to select the data to be
20 transferred from the first volatile cache memory to the first non-volatile storage media responsively to the mirror.